

Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) EP 0 959 601 A1

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
24.11.1999 Bulletin 1999/47

(51) Int Cl.<sup>6</sup>: H04L 29/06, H04L 12/56

(21) Application number: 99303688.8

(22) Date of filing: 12.05.1999

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(72) Inventor: Chauhan, Abhishek  
Palo Alto, California 94301 (US)

(74) Representative: Browne, Robin Forsythe, Dr.  
Urquhart-Dykes & Lord  
Tower House  
Merrion Way  
Leeds LS2 8PA West Yorkshire (GB)

(30) Priority: 21.05.1998 US 83864

(71) Applicant: SUN MICROSYSTEMS, INC.  
Palo Alto, California 94303 (US)

### (54) A system and method for server selection for mirrored sites

(57) The present invention is a system and method for providing server selection for mirrored sites. In a preferred embodiment, a user requests an address via a local name service, and the local name service requests the address from a server, examples of which include an authoritative name server, and a Optimizer Name Server (ONS). If the ONS has already calculated a best route which applies to this user, and a predetermined

time applying to the best route has not yet expired, then the ONS replies to the user with the best route server. If, however, the best route has not been determined or the predetermined time for the best route has expired, then the ONS replies to the user via a selection scheme, such as a round robin scheme, and also initiates a best route determination to determine which is the best route for this user.




FIG. 4

EP 0 959 601 A1

## Description

### BACKGROUND OF THE INVENTION

#### 1. Field of the Invention

[0001] The present invention relates to accessing data on the Internet via a computer system. In particular, the present invention relates to a system and method for selecting a server for requested data.

#### 2. Description of the Related Art

[0002] With the growing popularity of the Internet, providing speedy access to a requested location on the Internet is fast becoming a major issue. When a popular web page is hosted by a single computer, the Internet traffic to the computer can be overwhelming. To solve this problem, several computers can be utilized to host the web page such that each hosting computer, typically referred to as a server, contains a copy of the web page. If there are many servers at the same location, then the network connection to that location can become choked during a time of high Internet traffic. To avoid choking the network connection, mirrored servers can be located at different sites. These sites are herein referred to as mirrored sites. The use of multiple computers to host a web page is typically referred to as mirroring servers.

[0003] Figs. 1 and 2 illustrate basic communication lines in the Internet. Fig. 1 is a diagram showing the basic organization of the Internet. Groups of networks 100A-100C, 100D-100F, and 100G-100I, are shown to be connected via routers 102A-102F. A set of networks and routers combine to create an autonomous system (AS). A collection of networks within an AS can consistently perform various tasks with standards and protocols within the AS 104A-104C. Each AS 104A-104C is shown to be coupled with another AS via a gateway 106A-106B. Gateways 106A-106B typically contain enough information about the Internet such that each gateway has data related to every other gateway. Each gateway typically has routing tables which contains information relating to a route from point A to point B for the entire Internet.

[0004] Fig. 2 is a block diagram illustrating a path taken by a user's request for a particular address, and the path taken in receiving that address. Fig. 2 shows user 200 requesting an address, such as www.sun.com, from a local name service 201. The user's local name service 201 may be a service such as Netcom or AT&T. The local name service 201 then requests the address for sun.com from a domain name service (DNS) server 202a. Typically, the DNS server 202a requests the address for sun.com from another DNS server 202d, which in turn will typically request another DNS server 202c. This series of requests continues until the request for the address of sun.com arrives at an authoritative name service (Authoritative NS) 204. The authoritative NS 204

can determine the address of sun.com. In the example shown in Fig. 2, mirrored servers IP12 06a, IP2 206b, and IP3 206c are shown to be utilized. The authoritative NS 204 determines which server 206a-206c address to respond to the request for an address such as sun.com. The mirrored servers 206a-206c can be located in various geographically diverse locations such as an east location, west location, and an European location. An advantage of having mirrored servers located at various locations is that multiple locations can accommodate more Internet traffic than a single location. A further advantage of having mirrored servers in various locations is that if the user is located in Europe, accessing the European server 206c is most likely faster than the European user accessing the west coast server 206c.

[0005] The authoritative NS 204 can select a mirrored server based on several factors. These factors include the load on each server or the user's location. Ideally, the requested address should be returned to the user very quickly and download into the user's computer system as quickly as possible. The time between the request being sent out and the address being received is referred to latency. The time it takes to download the requested data is typically determined by the bandwidth.

[0006] Once the authoritative NS 204 determines which mirrored server the user should access, the authoritative NS 204 returns the address of the selected mirrored server back to the user via the route taken by the user's request.

[0007] A problem arises when a user bookmarks a particular web page. When a user bookmarks a particular web page, the user is typically bookmarking a particular server. For instance, if a user in California is sent to the European server 206c, and the user bookmarks the page, then when the user recalls the page from a bookmark, then the user will return to the European server 206c. However, the second time the user accesses the European server 206c, the European server may be heavily loaded, while the eastern server 206a may have a light load. In that situation, it may be better for the user to access the eastern server 206a rather than the European server 206c. Since the user has bookmarked a particular server, whatever scheme was instituted in the authoritative NS 204 is typically defeated. Additionally, if the user has bookmarked a particular server, it is better for the user if the bookmarked server can be accessed quickly. For instance, the user in California may be in a more optimal situation if the bookmarked server was located closer to the user than the European server 206c. These potential problems related to bookmarking typically occur in Http redirect mechanisms which are not transparent to the user.

[0008] The address given by the authoritative NS 204 is typically valid for a predetermined limited amount of time. If the user requests access to the same web page which the user had previously requested, the user should ask the authoritative NS 204 again for the address of the web page if the previous amount of time

has expired since the last request. In this manner, the authoritative NS 204 can reevaluate the situation and redetermine which server to route the user.

[0009] Several schemes for determining which of the mirrored servers to use has been utilized by various authoritative NS's 204. One such scheme is a load balancing named (LB named) system scheme. In this scheme, the authoritative NS 204 determines how many users are being served by each mirrored server and routes a new user to the mirrored server with the least load. The load is determined by the number of people accessing the server at a given time and the power of the server. The problem with the LB-named scheme is that although it factors load balancing, the scheme is indifferent to the location of the user. The location of the user is typically a factor in determining latency. If the user is very far from the assigned server, then it will typically take longer to access the far server than if the user was very close to the server, even if the closer server's load is heavier than the farther server's load.

[0010] Another scheme which is conventionally used by the authoritative NS 204 is the round robin DNS. The round robin DNS simply takes turns on which server is to be used. The problem with the round robin DNS scheme is that it does not take into consideration the various loads of the mirrored servers at any given time. Additionally, the round robin DNS scheme also fails to take into consideration the location of the user. Accordingly, the two factors of latency and bandwidth are not typically considered in the round robin scheme.

[0011] Yet another scheme can be referred to as an AS hop count scheme. In the AS hop count scheme, autonomous systems located between the user and a mirrored server are counted and the server with the smallest number of autonomous systems located between the user and the server is selected for use by that user. The AS hop count scheme assumes that the factors of bandwidth and latency are approximately reflected in the number of autonomous systems located between the user and a server. However, in practice, autonomous systems can be various sizes and a small autonomous system typically cannot accommodate as much traffic as a large autonomous system. Accordingly, a route including a larger number of autonomous systems may actually be faster than a route with a smaller number of autonomous systems if a very small AS is included in the route with the smaller number of autonomous systems.

[0012] What is needed is a system and method for a server selection for mirrored sites which optimizes access and is transparent to the user. The present invention addresses such a need.

### **SUMMARY OF THE INVENTION**

[0013] The present invention is a system and method for providing server selection for mirrored sites. In a preferred embodiment, a user requests an address via a

local name service, and the local name service requests the address from a server, examples of which include an authoritative name server, and an Optimizer Name Server (ONS). If the ONS has already calculated a best route which applies to this user, and a predetermined time applying to the best route has not yet expired, then the ONS replies to the user with the best route server. If, however, the best route has not been determined or the predetermined time for the best route has expired, then the ONS replies to the user via a selection scheme, such as a round robin scheme, and also initiates a best route determination to determine which is the best route for this user.

[0014] In the preferred embodiment, the best route determination begins with requesting all the mirrored servers for a round trip time to the requesting local name server used by the user. Each mirrored server then sends a name query to the local name server (LNS). The LNS then would reply with an error message stating that no such name exists. Each mirrored server then determines its own round trip time for the name query and the error message and sends the round trip time to the ONS. The ONS receives the round trip times of all mirrored servers and determines which mirrored server has the shortest round trip time (best route). The best route determination is then stored for a predetermined time. If the same LNS requests an address from the ONS within that predetermined time, then the mirrored server with the best route for that particular LNS is utilized.

[0015] Alternatively, rather than each mirrored server sending a name query to the LNS, each mirrored server sends an Internet control message protocol (ICMP) echo packet to the LNS. When the echo packet is returned from the LNS, then a round trip time can be determined by each mirrored server.

[0016] A method according to an embodiment of the present invention for computer server selection in a computer network including a plurality of servers is presented. The method comprising determining whether a best route has been determined; selecting one of the plurality of servers as a best route server if the best route has been determined; and selecting one of the plurality of servers as a selection scheme server scheme if the best route has not been determined. The selection of the selection scheme server is determined via a selection scheme, wherein the selection scheme does not include determining the best route.

[0017] A system according to an embodiment of the present invention for computer server selection in a computer network is also presented. The system comprises a plurality of servers; and an authoritative name server coupled to the plurality of servers. The authoritative name server determines whether a best route has been determined. It selects one of the plurality of servers as a best route server if the best route has been determined, and selects one of the plurality of servers as a selection scheme server if the best route has not been determined. The selection scheme does not include de-

termining the best route.

#### **Brief Description Of The Drawings**

[0018] Fig. 1 is a diagram illustrating the basic organization of the Internet.

[0019] Fig. 2 is a block diagram illustrating a path taken by a user's request for a particular address, and the path taken in receiving that address.

[0020] Fig. 3 is a block diagram of a computer system in which an embodiment of the present invention can reside.

[0021] Fig. 4 is a system diagram illustrating the data flow of the system for providing server selection for mirrored sites according to an embodiment of the present invention.

[0022] Fig. 5 is a flow diagram of a method according to an embodiment of the present invention for providing server selection for mirrored sites.

[0023] Fig. 6 is a flow diagram of a method according to an embodiment of the present invention for determination of best route as recited in step 508 of Fig. 5.

[0024] Fig. 7 is an alternative method according to an embodiment of the present invention for determination of best route recited in step 508 of Fig. 5.

#### **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS**

[0025] The following description is presented to enable one of ordinary skill in the art to make and to use the invention and is provided in the context of a patent application and its requirements. Various modifications to the preferred embodiments will be readily apparent to those skilled in the art and the generic principles herein may be applied to other embodiments. Thus, the present invention is not intended to be limited to the embodiment shown but is to be accorded the widest scope consistent with the principles and features described herein.

[0026] Figure 3 is a block diagram of a general purpose computer system suitable for carrying out the processing in accordance with one embodiment of the present invention. Figure 3 illustrates one embodiment of a general purpose computer system. Other computer system architectures and configurations can be used for carrying out the processing of the present invention. Computer system, made up of various subsystems described below, includes at least one microprocessor subsystem (also referred to as a central processing unit, or CPU) 302. That is, CPU 302 can be implemented by a single-chip processor or by multiple processors. CPU 302 is a general purpose digital processor which controls the operation of the computer system. Using instructions retrieved from memory 310, the CPU 302 controls the reception and manipulation of input data, and the output and display of data on output devices.

[0027] CPU 302 is coupled bi-directionally with mem-

ory 310 which can include a first primary storage, typically a random access memory (RAM), and a second primary storage area, typically a read-only memory (ROM). As is well known in the art, primary storage can be used as a general storage area and as scratch-pad memory, and can also be used to store input data and processed data. It can also store programming instructions and data, in the form of data objects and text objects, in addition to other data and instructions for processes operating on CPU 302. Also as well known in the art, primary storage typically includes basic operating instructions, program code, data and objects used by the CPU 302 to perform its functions. Primary storage devices 310 may include any suitable computer-readable storage media, described below, depending on whether, for example, data access needs to be bi-directional or uni-directional. CPU 302 can also directly and very rapidly retrieve and store frequently needed data in a cache memory (not shown).

[0028] A removable mass storage device 312 provides additional data storage capacity for the computer system, and is coupled either bi-directionally or uni-directionally to CPU 302. For example, a specific removable mass storage device commonly known as a CD-ROM typically passes data uni-directionally to the CPU 302, whereas a floppy disk can pass data bi-directionally to the CPU 302. Storage 312 may also include computer-readable media such as magnetic tape, flash memory, signals embodied on a carrier wave, PC-CARDS, portable mass storage devices, holographic storage devices, and other storage devices. A fixed mass storage 320 can also provide additional data storage capacity. The most common example of mass storage 320 is a hard disk drive. Mass storage 312, 320 generally store additional programming instructions, data, and the like that typically are not in active use by the CPU 302. It will be appreciated that the information retained within mass storage 312, 320 may be incorporated, if needed, in standard fashion as part of primary storage 310 (e.g. RAM) as virtual memory.

[0029] In addition to providing CPU 302 access to storage subsystems, bus 314 can be used to provide access other subsystems and devices as well. In the described embodiment, these can include a display monitor 318, a network interface 316, a keyboard 304, and a pointing device 306, as well as an auxiliary input/output device interface, a sound card, speakers, and other subsystems as needed. The pointing device 306 may be a mouse, stylus, track ball, or tablet, and is useful for interacting with a graphical user interface.

[0030] The network interface 316 allows CPU 302 to be coupled to another computer, computer network, or telecommunications network using a network connection as shown. Through the network interface 316, it is contemplated that the CPU 302 might receive information, e.g., data objects or program instructions, from another network, or might output information to another network in the course of performing the above-de-

scribed method steps. Information, often represented as a sequence of instructions to be executed on a CPU, may be received from and outputted to another network, for example, in the form of a computer data signal embodied in a carrier wave. An interface card or similar device and appropriate software implemented by CPU 302 can be used to connect the computer system to an external network and transfer data according to standard protocols. That is, method embodiments of the present invention may execute solely upon CPU 302, or may be performed across a network such as the Internet, intranet networks, or local area networks, in conjunction with a remote CPU that shares a portion of the processing. Additional mass storage devices (not shown) may also be connected to CPU 302 through network interface 316.

**[0031]** An auxiliary I/O device interface (not shown) can be used in conjunction with computer system. The auxiliary I/O device interface can include general and customized interfaces that allow the CPU 302 to send and, more typically, receive data from other devices such as microphones, touch-sensitive displays, transducer card readers, tape readers, voice or handwriting recognizers, biometrics readers, cameras, portable mass storage devices, and other computers.

**[0032]** In addition, embodiments of the present invention further relate to computer storage products with a computer readable medium that contain program code for performing various computer-implemented operations. The computer-readable medium is any data storage device that can store data which can thereafter be read by a computer system. The media and program code may be those specially designed and constructed for the purposes of the present invention, or they may be of the kind well known to those of ordinary skill in the computer software arts. Examples of computer-readable media include, but are not limited to, all the media mentioned above: magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and specially configured hardware devices such as application-specific integrated circuits (ASICs), programmable logic devices (PLDs), and ROM and RAM devices. The computer-readable medium can also be distributed as a data signal embodied in a carrier wave over a network of coupled computer systems so that the computer-readable code is stored and executed in a distributed fashion. Examples of program code include both machine code, as produced, for example, by a compiler, or files containing higher level code that may be executed using an interpreter.

**[0033]** It will be appreciated by those skilled in the art that the above described hardware and software elements are of standard design and construction. Other computer systems suitable for use with the invention may include additional or fewer subsystems. In addition, bus 314 is illustrative of any interconnection scheme serving to link the subsystems. The computer system

shown in Fig. 3 is but an example of a computer system suitable for use with the invention. Other computer architectures having different configurations of subsystems may also be utilized.

**[0034]** Fig. 4 is a block diagram of a system according to an embodiment of the present invention for providing service selection for mirrored sites. Fig. 4 also shows the data flow between components. A user 400 sends a request for an address to the user's local name service 402. The local name service 402 then requests the address from an Optimizer Name Server (ONS) 404. The request from the local name server (local NS) 402 may route through several servers prior to reaching the ONS 404.

**[0035]** The ONS 404 is preferably a name server which can specifically answer queries, such as name queries. Although the ONS can work with more than one name, it is limited to approximately two or three names. Alternatively, the name server 404 can be a conventional ANS. Accordingly, a conventional ANS can be used either in addition to a ONS or in place of a ONS to handle a greater number of names. The ONS and the mirrored servers are preferably implemented in JAVA™. Further details of the functions of the ONS will be described in conjunction with Figs. 5, 6, and 7.

**[0036]** The data flow shown in Fig. 4 can best be understood when referred to in conjunction with Figs. 5, 6, and 7. Fig. 5 is a flow diagram of a method according to an embodiment of the present invention for server selection for mirrored sites. Fig. 6 is a flow diagram of a method for best route determination of step 508 in Fig. 5. Fig. 7 is an alternative method for the best route determination of step 508 in Fig. 5.

**[0037]** The method according to the present invention for providing service selection for mirrored sites can be seen in the examples shown in Figs. 4 and 5. The user 400 requests an address with a name to the local NS (LNS) 402 in step 500. The LNS eventually contacts and requests the address from the ONS 404 via step 502. It is then determined by the ONS if a best route has already been determined for this LNS which is still valid via step 504. A best route determination is preferably only valid for a predetermined time, such as thirty minutes. If a best route has been determined for this LNS which is still within the predetermined time, then the ONS 404 replies to the LNS 402 with the best route mirrored server via step 506.

**[0038]** If, however, the best route has not been determined for this LNS, or the predetermined time has expired, then the ONS 404 replies to LNS 402 via a selection scheme, such as a round robin scheme, via step 508. Additionally, if a best route has not been determined for this LNS, a best route determination for the LNS is initiated with step 508.

**[0039]** A preferred method according to an embodiment of the present invention of best route determination (step 508 of Fig. 5) can be described in connection with references to Figs. 4 and 6. In initiating best route

determination, the ONS 404 requests from all the mirrored servers 406a-406b the round trip time to the LNS 402 via step 600. Each mirrored server 406a-406b then sends a name query to the LNS 402 via step 602. The LNS 402 then replies with an error message stating that the requested name does not exist via step 604. Each mirrored server 406a-406b then determines the round trip time from the time the name query was sent to the time the error message was received via step 606. Each mirrored server 406a and 406b then sends its round trip time to the ONS 404 via step 608. The ONS 404 can then determine which mirrored server has the best route for this particular LNS 402 via step 610. The best route is equivalent to the fastest route. The best route is then stored for a predetermined time via step 612.

[0040] When the LNS 402 requests an address via step 614, it is determined whether the pre determined time has expired in step 616. If so, then the ONS 404 replies to the LNS 402 via a selection scheme, such as a round robin scheme, via step 620 and a best route determination is initiated via step 600. If, however, the predetermined time has not expired via step 616, then the LNS 402 is assigned to the mirrored server 406 with the best route via step 618.

[0041] An alternative method according to an embodiment of the present invention for the best route determination (step 508 of Fig. 5) can be described in conjunction with Figs. 4 and 7. The ONS 404 requests from all the mirrored servers 406a-406b the round trip time to the LNS 402 via step 700. Each mirrored server 406a-406b then sends an echo packet to LNS 402 via step 702. An echo packet is an Internet Control Method Protocol (ICMP) packet which is returned by the receiver when the receiver receives the echo packet. Echo packets are commonly known in the industry. Each mirrored server 406a-406b then receives a return echo packet from the LNS 402 via step 704. Each mirrored server 406a-406b then determines a round trip time and sends the round trip time to the ONS 404 via step 706. The remaining step 708-718 are the same as steps 610-620 of Fig. 6.

[0042] A method and system for providing server selection for mirrored sites has been disclosed. Software written according to the present invention may be stored in some form of computer readable medium such as memory or CD-ROM or transmitted over a network, and executed by a processor.

[0043] Although the present invention has been described in accordance with the embodiments shown, one of ordinary skill in the art will readily recognize that there could be variations to the embodiment and these variations would be within the spirit and scope of the present invention. Accordingly, many modifications may be made by one of ordinary skill in the art without departing from the spirit and scope of the appended claims.

## Claims

1. A method for computer server selection in a computer network including a plurality of servers, the method comprising:

determining whether a best route has been determined;  
selecting one of the plurality of servers as a best route server if the best route has been determined; and  
selecting one of the plurality of servers as a selection scheme server if the best route has not been determined, the selection of the selection scheme server being determined via a selection scheme, wherein the selection scheme does not include determining the best route.

2. The method of claim 1, further including:

determining if a predetermined time has expired if the best route has been determined; and  
selecting a selection scheme server by using a selection scheme if the predetermined time has expired.

3. The method of claim 1 or 2, wherein the selection scheme is a round robin scheme.

4. The method of any of the preceding claims further including:

determining the best route if the best route has not been determined, wherein the determination of best route includes determining a round trip time from one of the plurality of servers to a local name server.

5. The method of claim 4, wherein the determination of the round trip time includes sending a name query to the local name server and receiving message from the local name server.

6. The method of claim 4, wherein the determination of the round trip time includes sending an echo packet to the local name server and receiving a returning echo packet from the local name server.

7. The method of any of the preceding claims, wherein the determination of best route includes determining a first round trip time from a first one of the plurality of servers to a local name server and a second round trip time from a second one of the plurality of servers to the local name server.

8. The method of claim 7, wherein the first round trip time and the second round trip time is compared to determine which one of the first and second servers

has the best route.

9. The method of any of the preceding claims, wherein the best route includes a shortest round trip time between one of the plurality of servers and a local name server. 5

10. A system for computer server selection in a computer network including a plurality of servers, the system comprising: 10

means for determining whether a best route has been determined;

means for selecting one of the plurality of servers as a best route server if the best route has been determined; and 15

means for selecting one of the plurality of servers as a selection scheme server if the best route has not been determined, the selection of the selection scheme server being determined via a selection scheme, wherein the selection scheme does not include determining the best route. 20

11. A system for computer server selection in a computer network, the system comprising: 25

a plurality of servers; and

an authoritative name server coupled to the plurality of servers, wherein the authoritative name server determines whether a best route has 30

been determined, and selects one of the plurality of servers as a best route server if the best route has been determined; and selects one of the plurality of servers as a selection scheme 35

server, if the best route has not been determined, wherein the selection scheme does not include determining the best route.

12. A computer program product for computer server selection in a computer network including a plurality of servers, comprising: 40

computer code determining whether a best route has been determined; 45

computer code selecting one of the plurality of servers as a best route server if the best route has been determined;

computer code selecting one of the plurality of servers as a selection scheme server if the best route has not been determined, the selection of the selection scheme server being determined via a selection scheme, wherein the selection scheme does not include determining the best route; and 50

a computer readable medium that stores the computer codes. 55




FIG. 1



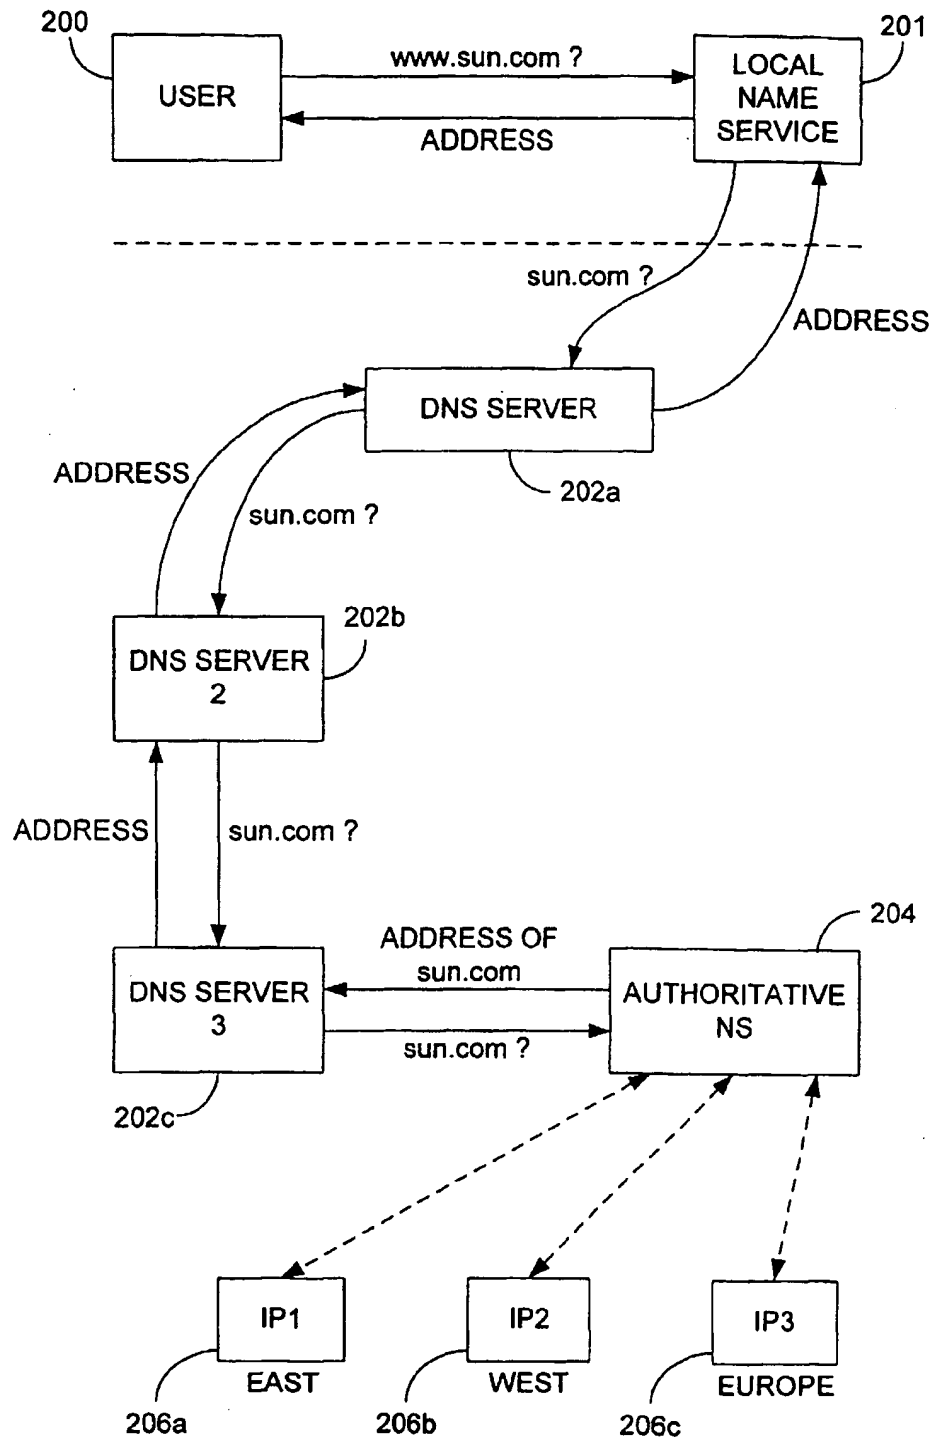


FIG. 2




FIG. 3




FIG. 4

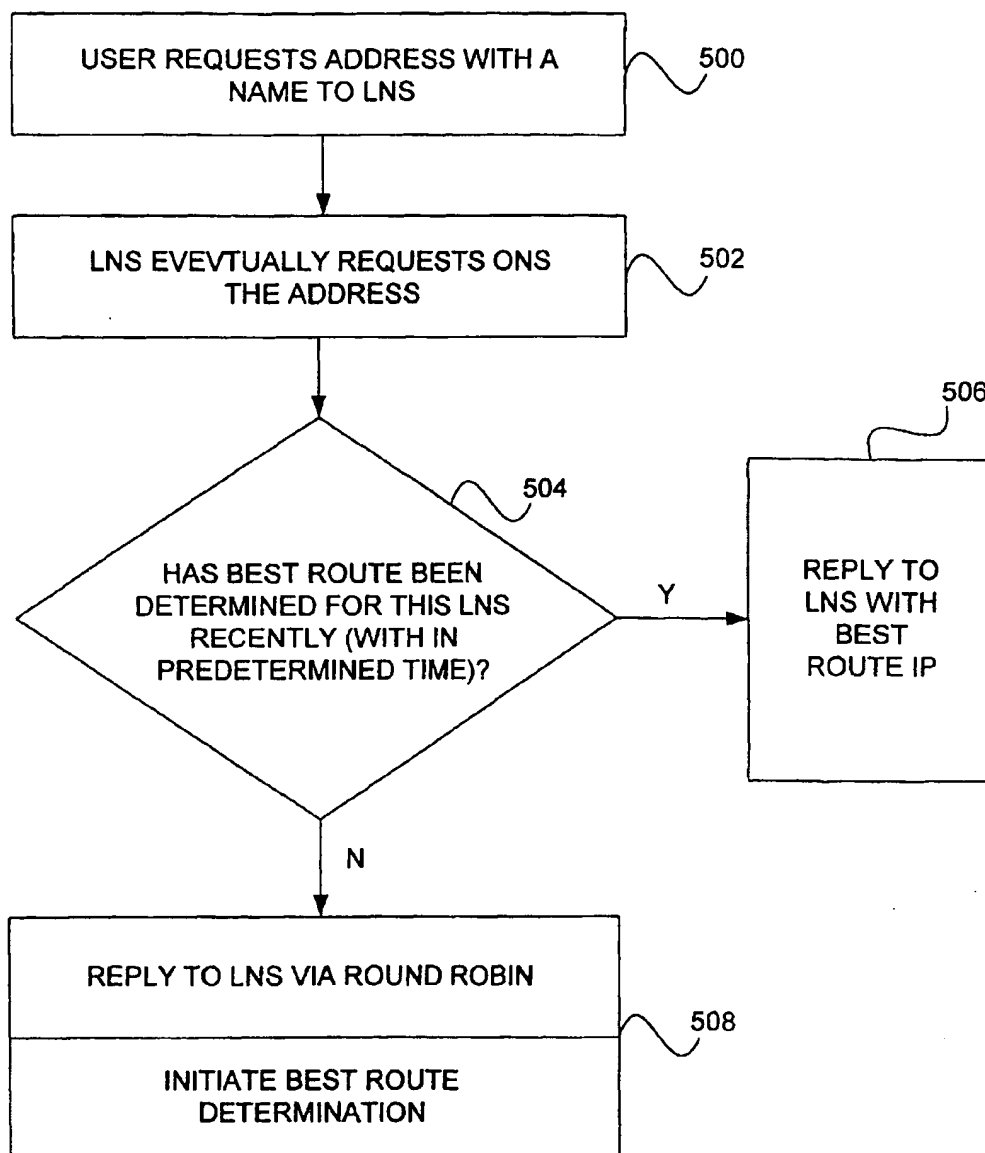


FIG. 5

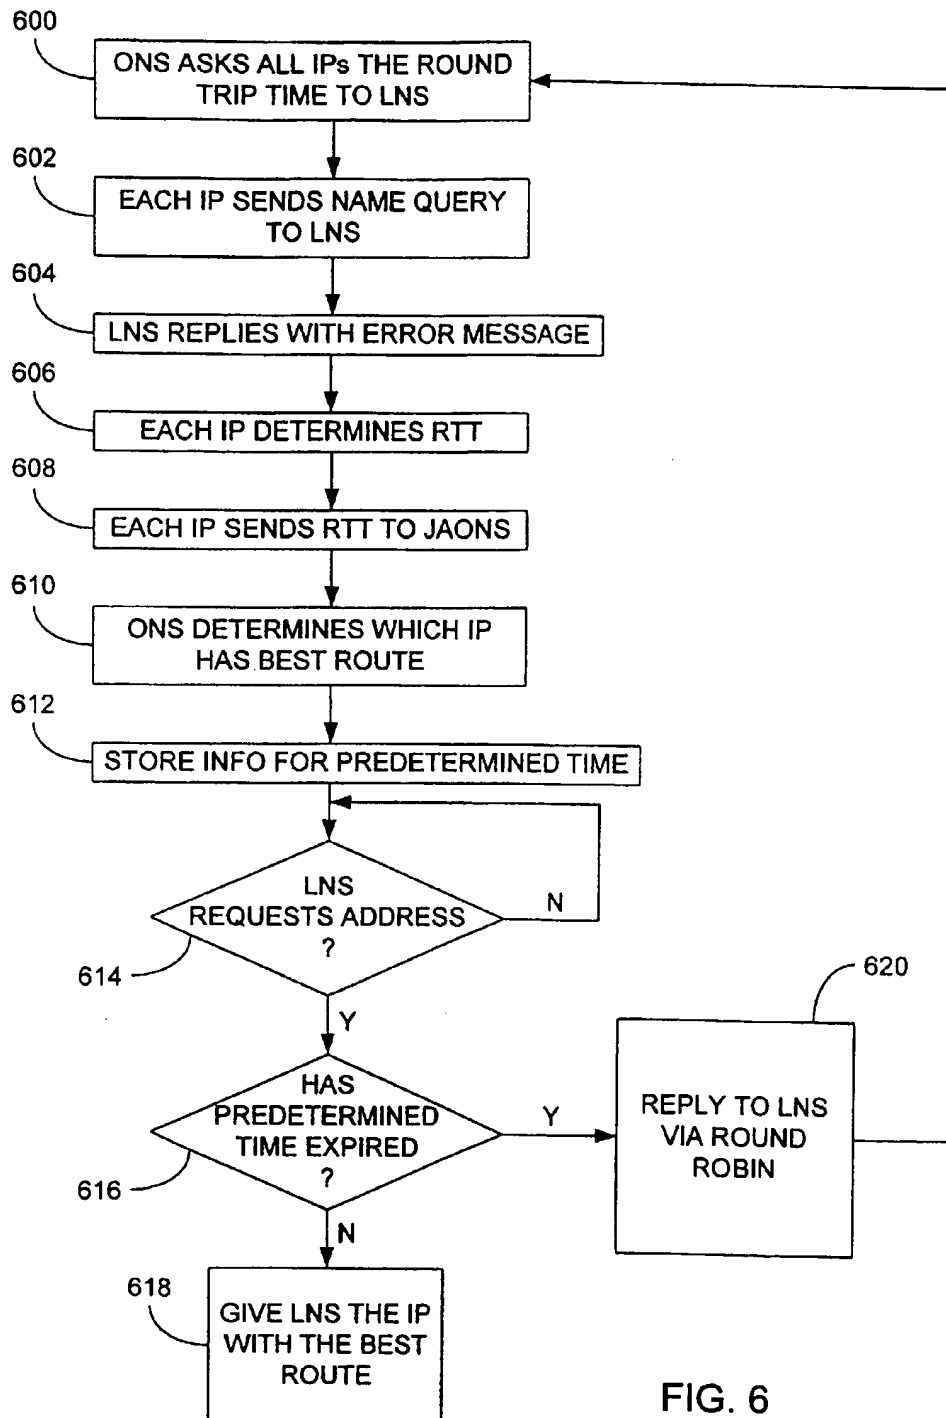


FIG. 6

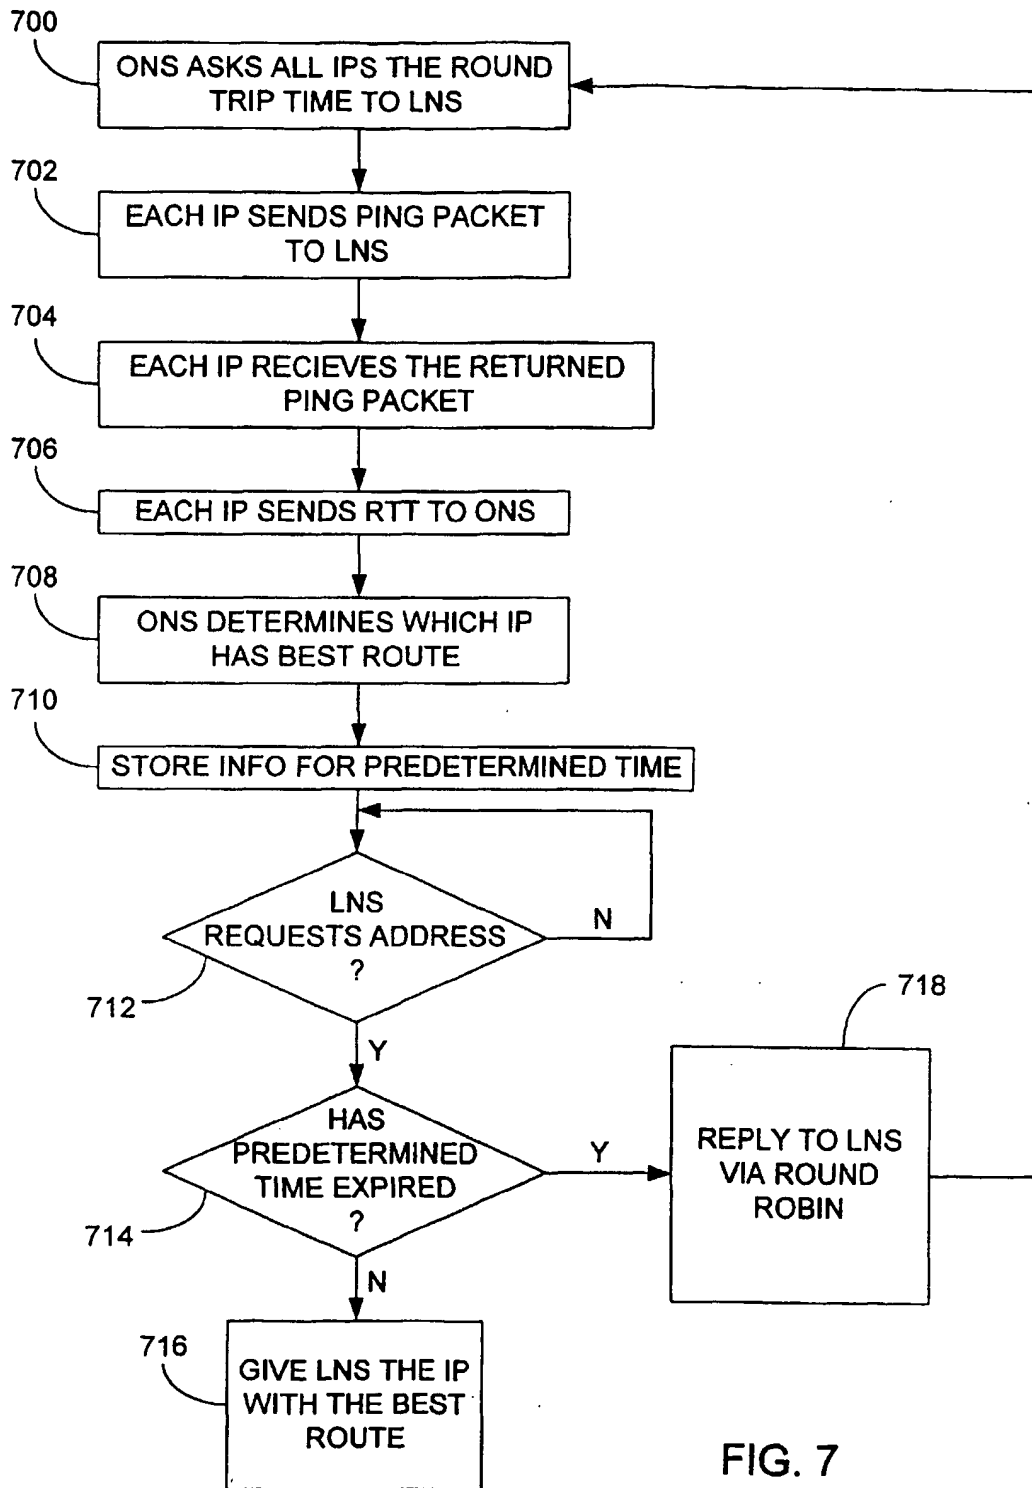


FIG. 7



European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number  
EP 99 30 3688

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	BHATTACHARJEE S ET AL: "Application-layer anycasting" PROCEEDINGS OF THE IEEE INFOCOM '97. THE CONFERENCE ON COMPUTER COMMUNICATIONS. 16TH ANNUAL JOINT CONFERENCE OF THE IEEE COMPUTER AND COMMUNICATIONS SOCIETIES. DRIVING THE INFORMATION REVOLUTION, KOBE, APRIL 7 - 12, 1997, vol. 3, 7 April 1997 (1997-04-07), pages 1388-1396 1396 vol.3, XP002109465 INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS ISBN: 0-8186-7782-1	1,10-12	H04L29/06 H04L12/56
Y	* page 1388, left-hand column, paragraph 2	2,3	
	- page 1394, left-hand column, paragraph 1		
A	*	4-9	
Y	COLAJANNI M ET AL: "ADAPTIVE TTL SCHEMES FOR LOAD BALANCING OF DISTRIBUTED WEB SERVERS" PERFORMANCE EVALUATION REVIEW, vol. 25, no. 2, 1 September 1997 (1997-09-01), pages 36-42, XP000199853 ISSN: 0163-5999 * page 36, paragraph 4 *	2	TECHNICAL FIELDS SEARCHED (Int.Cl.6) H04L G06F
Y	COLAJANNI M ET AL: "SCHEDULING ALGORITHMS FOR DISTRIBUTED WEB SERVERS" PROCEEDINGS OF THE 17TH. INTERNATIONAL CONFERENCE ON DISTRIBUTED COMPUTING SYSTEMS, BALTIMORE, MD., MAY 27 - 30, 1997, 27 May 1997 (1997-05-27), pages 169-176, XP000793033 INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS ISBN: 0-8186-7814-3 * page 169, right-hand column, paragraph 2 - paragraph 3 *	3	
		-/-	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 27 September 1999	Examiner RAMIREZ DE AREL., F
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 03 82 (P04001)



European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 99 30 3688

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.5)
A	WO 98 18076 A (INTERVU INC) 30 April 1998 (1998-04-30) * page 17, line 15 - page 18, line 37 *	1,4-9	
P,X	WO 98 31107 A (GIFFORD DAVID K) 16 July 1998 (1998-07-16) * page 14, line 10 - page 16, line 33 *	1,2,7-12	
			TECHNICAL FIELDS SEARCHED (Int.Cl.5)
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 27 September 1999	Examiner RAMIREZ DE AREL..., F.
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 03 82 (104/001)



ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.

EP 99 30 3688

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on  
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

27-09-1999

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9818076 A	30-04-1998	AU 5152298 A EP 0932866 A	15-05-1998 04-08-1999
WO 9831107 A	16-07-1998	NONE	

EPO FORM P436

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82